arXiv:1611.08841v1 [cs.CV] 27 Nov 2016

Long-Term Image Boundary Extrapolation

Apratim Bhattacharyya, Mateusz Malinowski, Bernt Schiele, Mario Fritz
Max Planck Institute for Informatics, Saarbrücken, Germany
abhattac , mmalinow, schiele, mfritz@mpi-inf.mpg.de

Figure 1: We predict the future of videos. (first row): the true future frame for $t+4$. (second row): RGB extrapolation of [15] which gets blurry quickly. (third row): our boundary extrapolation overlayed on the unobserved, true future frame, which outlines the boundaries precisely and enables long term prediction.

Abstract

Boundary prediction in images and videos has been a very active topic of research and organizing visual information into boundaries and segments is believed to be a corner stone of visual perception. While prior work has focused on predicting boundaries for observed frames, our work aims at predicting boundaries of future unobserved frames. This requires our model to learn about the fate of boundaries and extrapolate motion patterns. We experiment on established real-world video segmentation dataset, which provides a testbed for this new task. We show for the first time spatio-temporal boundary extrapolation, that in contrast to prior work on RGB extrapolation maintains a crisp result. Furthermore, we show long-term prediction of boundaries in situations where the motion is governed by the laws of physics. We argue that our model has with minimalistic model assumptions derived a notion of "intuitive physics".

1. Introduction

Humans possess the skill to imagine future states of on observed scene, which supports tasks ranging from planning to object manipulation. Observing a moving ball, we have a reasonably good estimate about the future trajectory of the ball, which e.g. is key for the goal keeper to catch the ball.

This has inspired a recent line of research that aims at capturing this type of intuitive understanding of scene physics via machine learning techniques – which is often called "intuitive physics". E.g. [5] predicts futures states of balls moving on a billiard table and [12, 13] predict the stability and future states of towers made out of blocks. These models typically are parametric of some sort, are learnt from data or only predict a qualitative outcome of the scene. Moreover, both [5] and [12] have an "object notion", meaning that the model knows a priori the location or type of the objects it is supposed to model.

Recently, full future frame predication of natural (RGB) videos has been studied that is agnostic to the underlying

objects and causes of the change depicted in the sequence [19, 15]. But up to now, only very short range predictions of a few frames have been shown and blurriness/distortion artifacts occur in the predicted future frames as shown in Figure 1(second row) and also later evaluated in Figure 4a.

On the other hand, a lot of progress has been made in the field of image and video segmentation supported by datasets like VSB100 [7]. The segmented natural videos discard many details of the full RGB videos mentioned above which are hard for a model to learn and yet these segments still capture the important structures and extents of objects. The boundaries between segments gives rise to boundary images. The gray-scale boundary image encodes boundary location information. In many scenarios this high frequency information is crucial for meaningful predictions about the future, e.g. on a billiard table the location of ball and table boundaries in the past are necessary to infer the state of the table (location of the balls) in the future. Figure 1(last row) shows example results of our method that accurately extrapolates such image boundaries into the future (overlayed with groundtruth for reference), while also not being subject to the before mentioned issues blur.

Our main contribution is the first model to extrapolate future boundary frames of segmented videos and to explore the performance of these models under structured motion – deterministic and non-deterministic ones. We evaluate performance both on natural video sequences from VSB100 and synthetic and real billiard table sequences. We demonstrate that the boundaries predicted by our models are accurate over long time horizons (Figure 4b shows the mean squared error of predicted boundaries vs RGB) and can be used for sharper RGB frame prediction. Moreover, we show that our models can develop an intuitive understanding of physics from raw visual input without any strong parametric model of the motion or "object notion" and is capable of making accurate long term predictions under deterministic settings.

2. Related work

RGB frame prediction. This problem has been recently explored in [21, 19, 15]. The work of [21] focused on learning representations of video sequences. They used an LSTM encoder unit to encode videos into a vector which they used for predicting future frames. However, they focused on low-resolution input sequences and predicted frames has blurriness problems. In [19], the authors focused on the problem of blurring caused by using the mean squared loss as an objective function. They sought to remedy this problem by discretizing the input through k-means atoms and predicting on this vocabulary instead. The work of [15] also focused on this problem. They proposed using adversarial loss, which lead to improved results over [19]. These approaches produce sharp short term predictions (1 or 2 frames in the future) but still suffer from blurriness problems in the long term already starting at 3 frames into the future. Moreover, these works have focused on natural videos or datasets like MNIST digits. The ability of these models to learn the dynamics of deterministic motion has not been explored. Works of [22, 16] predict frames of videos of bouncing balls, but their dataset is very limited in size and resolution. Moreover, they do not consider generalization with respect to the number of balls and their velocities.

Fitting physics models to video. Works like that of [23, 17] focus on predicting outcomes of physical events in videos or images. In [23] the authors propose "Galileo" that estimates the physical properties of objects and inverts a physics engine to predict outcomes. While in [17], the authors predict the motion of objects using a single query image using a neural network which matches the image to a moment in a video which is closest in describing the dynamics of motion of the scene depicted in the image. A key difference to our work is, that we are not relying on a "object notion".

Intuitive physics. Developing an intuitive understanding of physics from raw visual input has been recently explored in [5, 12, 13]. [5] predicts future states of balls moving on a billiard table. [12, 13] predict the stability of towers made out of blocks. The work of [12] can also predict future locations of the blocks. However, both [5] and [12] have an "object notion". [13] focuses only on predicting the outcome not the exact state.

Video segmentation. Video segmentation as the task of finding consistent spatio-temporal boundaries in a video volume has received significant attention over the last years [6, 18, 7, 2], as it provides an initial analysis and abstraction for further processing. In contrast, our approach aims at extrapolating these boundaries into the future without any video observed for the future frames.

3. Models

We base our models on the recent success of Deep Learning. We approach long-term extrapolation by recursive schemes, which allows for efficient long term extrapolation. However, this means that errors are potentially propagated and accumulated over time. In order to mitigate such effects, we need very accurate models that consolidate information. By analyzing prior work on frame prediction (see section 2), we identify several key challenges such models have to address.

Large Spatio-Temporal Receptive Field. The output layer neurons should have a wide receptive field [8] to preserve long range spatial and temporal dependencies and learn about interaction with other boundaries in a spatio-temporal context.

Preserving Resolution and Preventing Blurred Output. The models must maintain resolution in order to derive

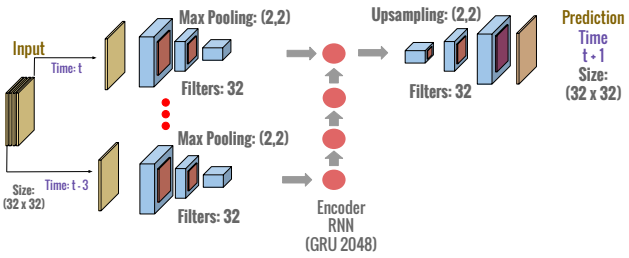

(a) Convolutional RNN - encoder decoder architecture

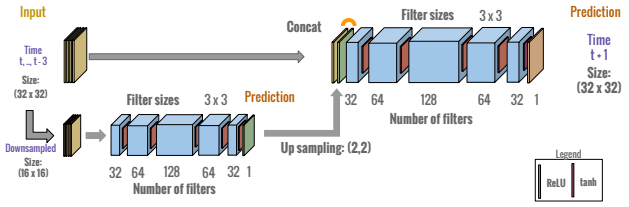

(b) Convolutional Multi-scale architecture (only 2 out of 4 scales illustrated)

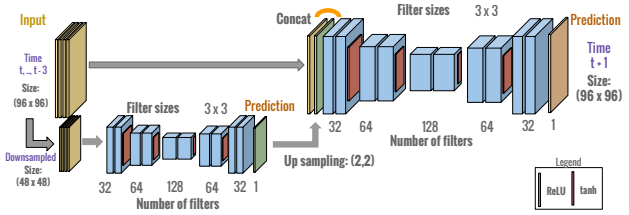

(c) Convolutional Multi-scale with Context architecture (only 2 out of 4 scales illustrated)

Figure 2: Our model architectures.

a high fidelity output boundary map. Excessive pooling or tight bottlenecks with fully connected layers have been shown to be successful in classification tasks, but also have shown to induce image degradations for image synthesis tasks [19].

Globally Consistent Predictions from Local Models. In order to generalize across diverse input sequences while maintaining a tractable number of parameters, a patch based approach is advised. Here, the models observe and extrapolate on patches rather than the complete input image. However, for accurate long term prediction, it is crucial to ensure global consistency through communication between the patch extrapolators. Consider a video of a moving ball. The trajectory of a ball might intersect with multiple patches. To correctly extrapolate the motion far into the future, instances of a model predicting on neighboring patches need to be consistent especially during transition of the ball between patches.

We propose the following model architectures for boundary extrapolation to meet the key challenges described above. We build upon the models used for frame prediction, while adapting and improving them for long term boundary prediction.

3.1. Convolutional-RNN (C-RNN)

We extend the model in [21] with convolutional layers as shown in Figure 2a, which can extract high quality location invariant features and have been very successful in various tasks [9]. The model consists of an convolutional-encoder GRU unit which reads in the input frame sequence one time step at a time and produces a single vector as a summary. This encoder unit contains of three convolutional layers with pooling and ReLU non-linearities in between, which extracts features which are then read by the GRU unit. The convolutional layers have 32 filters each of size 3x3. The summary vector is read by a convolutional-decoder unit to produce the final output frame. The decoder unit consists of a dense layer followed by convolutional layers with upsampling in between to maintain resolution. Upsampling can be thought of as convolution with a fractional stride. The dense layer along multiple convolutional layers in the decoder unit creates a wide receptive for the output layer neurons. Thus, the model fulfills the first key challenge.

3.2. Convolutional Multi-Scale (CMS)

The previously discussed Convolutional-RNN can also be understood as a type of auto-encoder [21] due to the intermediate summary vector, which acts a tight bottleneck layer. Thus, it fails to meet our second key challenge. One way to deal with this and to meet the second key challenge is to use end to end convolutions instead. We consider a Multi-Scale model architecture akin to a Laplacian pyramid as shown in Figure 2b. Such a model architecture has been used for generating natural images [3] and predicting future natural frames [15] (also see section 2). It contains multiple levels which observes the input frame(s) at increasing (coarse to fine) scales. The input I to a certain level (L_{2k}) are the scaled input frame(s) X_{2k} and the extrapolated frame(s) O from the previous coarser level (L_k), which are upsampled $\hat{O}$ to the scale at the current level.

$$I(L_{2k}) = \left\{X_{2k}, \hat{O}(L_k)\right\}$$

The coarse extrapolated frames $\hat{O}(L_k)$ act a guide for each level of the model.

In detail, we use four levels, with scales increasing by a factor of two. Each level of the model consists of five convolutional layers with 32, 64, 128, 64 and 32 filters respectively of constant size 3x3. We use *ReLU* non-linearities between every layer and a *tanh* at the end (ensures output in range [0,1]). This creates a wide receptive field for the output layer neurons and meets the first key challenge.

However, the receptive fields are not uniform in size. The neurons at the boundary of the (2d) output layer have a smaller receptive field compared to the neurons at the center. This leads to a non-uniform (training and test) error distribution at the output layer neurons. In Figure 3 we plot

the average error at the output layer neurons at increasing distance from the patch border. Error increases consistently from patch center (right) to the patch border (left).

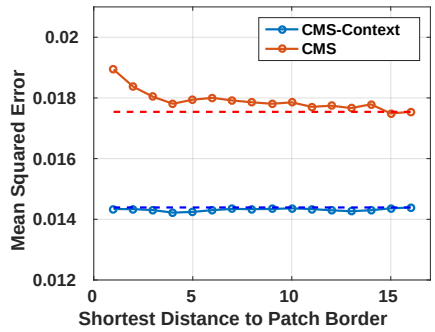

Figure 3: Mean squared error with increasing distance from the patch border. The Convolutional Multi-Scale model (red) has higher error near the patch boundary vs. Convolutional Multi-Scale Context (blue).

3.3. Convolutional Multi-Scale with Context (CMSC)

For the two models introduced previously, instances for the model extrapolating on different patches cannot communicate and therefore do not address the third challenge. To enable communication, we build upon the Convolutional Multi-Scale described before and introduce a context. This context model observes a central patch along with the directly neighbouring 8 patches – which we call the context. However, the model only extrapolates the central patch. While predicting recursively, the model observes its previous output along with the the output of the neighboring patch extrapolators. This allows for communication and constitutes a "read-write" architecture, enabling the learning for consistent prediction and hereby addressing our third challenge.

To deal with the larger input size, we make each level of the Muilt-Scale model deeper by doubling the number of convolutional layers and introduce a moderate amount of pooling (see Figure 2c). In addition, we double the number of filters after each pooling step (32, 64 and 128). Pooling is followed by upsampling layers to maintain resolution. This ensures large receptive fields at the output layer of each level without introducing a tight bottleneck layer. Hence, this model meets all three key challenges. Moreover, the addition context lead to uniform receptive fields and a uniform error distribution as cann be seen in Figure 3.

3.4. Loss

We use the L2 loss during training of the previously described models. In RGB frame prediction a variety of loss functions have been explored in combination with L2 loss e.g. Adverserial loss [15] to prevent blurring. This is because the probability distribution for an output pixel is multimodal and L2 loss predicts the mean. In contrast, our predicted boundary maps are essentially binary which encode boundary location information. Moreover, in a pilot study we could not observe improvements by adding an advererial loss for boundary prediction.

4. Experiments

We evaluate the models in section 3 both on natural video sequences and sequences with structured, deterministic motion. We convert each video into 32x32 pixel patches. The Convolutional-RNN and Multi-Scale models are trained on these patches. The Multi-Scale model with Context is trained on patches along with the eight neighbouring patches which results in a context of size 96x96 pixels. We use the ADAM optimizer during training. As we want sharp and accurate boundaries, we use the established boundary precision recall (BPR) evaluation metric from the video segmentation literature [7]. This metric is defined for a set P of predicted boundary images and G of corresponding ground truth boundary images as,

$$P = \frac{\sum_{B_p \in P, B_g \in G} \mid B_p \cap B_g \mid}{\sum_{B_p \in P} \mid B_p \mid}$$

$$R = \frac{\sum_{B_p \in P, B_g \in G} \mid B_p \cap B_g \mid}{\sum_{B_g \in G} \mid B_g \mid}$$

$$F = \frac{2PR}{P+R},$$

where P is boundary precision, R is boundary recall and F is the combined F-measure.

4.1. Short term extrapolation on natural video sequences

Dataset and Training. We use the VSB100 dataset which contains 101 videos with a maximum 121 frames each. We randomly choose 30 videos for training and the remaining 71 videos for testing from the VSB100 dataset. The videos contain a wide range of objects of different sizes and shapes. The videos also have a wide variety of both object and camera motion. We use the hierarchical video segmentation algorithm in [10] to segment these videos. The output is a ultra-metric contour map (ucm). Boundaries higher in the hierarchy typically correspond to semantically coherent entities like animals, vehicles etc and are therefore temporally more stable. We discard boundaries belonging to the lowest level of the hierarchy, as they are temporally very unstable. However, we keep the rest intact and use their values as a confidence measure on boundary location at a pixel.

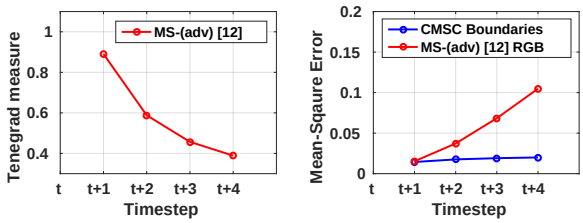

(a) Tenengrad measure. (b) Mean squared error.

Figure 4: Error analysis of RGB and boundary prediction with time.

Methods & baselines. We first evaluate and compare RGB prediction to boundary extrapolation on VSB100. Then, we evaluate the models in section 3 to extrapolate boundaries of segmented VSB100 videos. Recall that, the ground-truth boundaries (ucm) in VSB100 have different confidence values. Thus, we threshold the extrapolations before comparison to the groundtruth. We vary the threshold to obtain a precision-recall curve and report the area under the curve (AUC) along with the best F-measure across all thresholds. We include a "last frame" baseline by using the last input frame as constant extrapolation and a "Optical flow" baseline. As many boundaries do not change between frames in the videos of VSB100, the last input is not a bad baseline especially when we are predicting one step into the future. In case of the optic flow baseline, the optic flow is calculated between the last and the second last input frames (the frames at time t - 1 and t) using the Epic flow method of [20]. The boundary pixels at time t are propagated using the calculated flow to generate extrapolations at t + 1, t + 2 and t + 4. We also include the Multi-Scale model (MS - (adv)) from [15], which is trained on the segmented VSB100 videos with a combination of L2 and adversarial loss (with default parameters). We also evaluate the RGB predictions on VSB100 using the model of [15].

RGB verses boundary prediction. We report the sharpness of RGB frames (of VSB100) predicted by [15] using the Tenengrad measure [11] in Figure 4a. The Tenengrad measure pools the gradient magnitude information of the image. We observe that the model of [15] makes increasingly blurry predictions into the future, which is also confirmed the visual result in Figure 4. We also compare the mean squared error of RGB predictions of [15] and extrapolated boundaries of our Convolutional Multi-Scale Context model in Figure 4b. We are a sharper increase in the error of RGB predictions compared to boundaries in the long term.

Discussion of results on boundaries. We report the quantitative results results in Figure 5a and Figure 5b and the qualitative results in Figure 6.

Quantitative evaluation: Overall, our Convolutional Multi-Scale Context model (MS - Context, red lines) significantly outperforms the others. Our Convolutional Multi-Scale architecture is second. This demonstrates the importance of the large receptive fields and context (first and third key challenge in section 3). The good performance of both of the fully convolutional models versus the Convolutional - RNN model, shows that tight bottleneck layers cause degradation in performance (second key challenge section 3). Inferior performance of the Multi-Scale models [15] applied to boundaries shows that Adversarial loss is not well suited to the boundary extrapolation problem. The poor performance of the "Optic flow" baseline is due to not accurate flow information at object boundaries.

Qualitative evaluation: The boundaries produced by both of our Convolutional Multi-Scale architectures are sharp in case of deterministic and smooth motion e.g. the predictions in Figure 6 from the videos airplane and dominoes. However, the models are not able to deal with large or non-deterministic motion (e.g. of human actors). The models in such situations react by blurring the boundaries, as a consequence of using the mean squared error. While predicting recursively, this leads to loss of boundary confidence and eventual vanishing boundaries. Both the Convolutional - RNN and the Multi-Scale model of [15] produce blurry boundaries. The "Optic flow" baseline produces discontinuous (jagged) boundaries. (See section 6 for more examples).

Sharpening RGB predictions with fusion. The sharp boundaries produced by our models raises the prospect of sharpening RGB predictions in a fusion scheme. We show promising, initial results in Figure 7. We produce sharpened fused RGB predictions by taking the derivative of the RGB images and amplifying the gradients at the boundary locations (given by our extrapolated boundaries) in the spirit of [4]). This procedure improves the Tenengrad measure at t + 2 from 0.81 to 0.96 and at t + 4 from 0.49 to 0.63 of the Dominoes sequence.

4.2. Long term extrapolation on sequences with structured deterministic motion

Dynamics of motion in the videos in the VSB100 dataset is frequently very complex and involve difficult to predict and non-deterministic actions of actors. This convolves the challenges of long-term predictiong with a non-deterministic that cannot be captured by our model and hence would be treated as noise. Therefore we evaluate the performance for long-term prediction of the models on structured, deterministic motion. We evaluate on both real and synthetic billiard ball sequences. The motion of the balls in such sequences are deterministic and we aim to extrapolate them over long time-steps. We begin by describing how the sequences were collected.

Synthetic data generation. The synthetic billiard ball sequences are sampled from worlds which consists of balls moving on a frictionless surface with a boundary, akin to a billiard table. We used pygame to create such worlds and

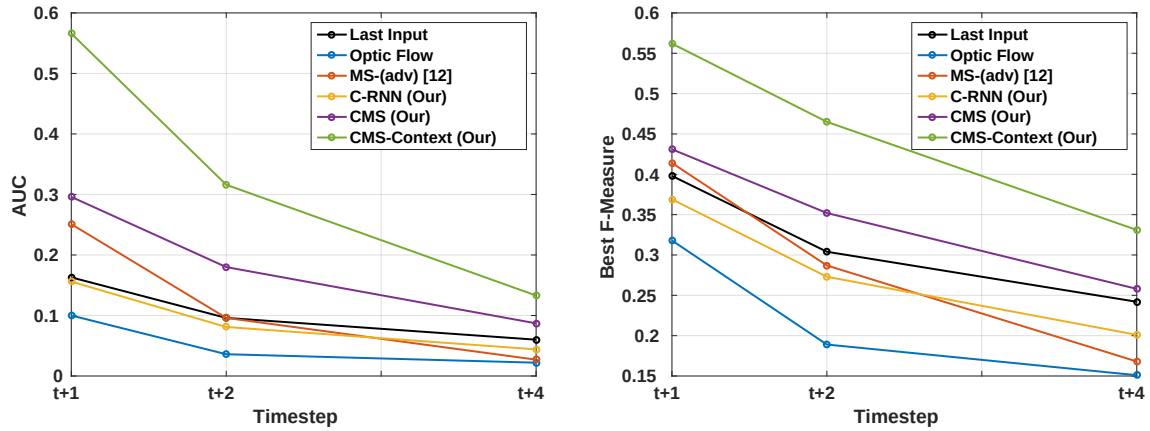

(a) Area under the curve. (b) Best F-measure.

Figure 5: Evaluation of the models in section 3 on VSB100.

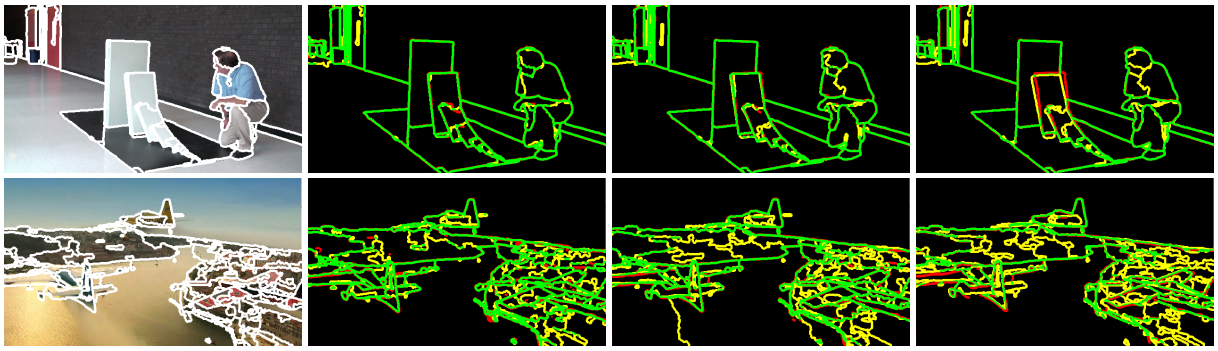

Last Input Frame Extrapolation: t + 1 Extrapolation: t + 2 Extrapolation: t + 4

Figure 6: Rows top to bottom: Extrapolations on dominoes and airplane sequences from VSB100. Correct boundaries predictions are encoded in green. Missed boundaries are encoded in yellow. Wrong boundaries are encoded in red.

sample boundary images from them. The output images contain boundaries that can stem from ball(s) or the table and are binary (indicating a boundary at that location). During evaluation, as the target is always a binary image, we report only the best F-measure obtained by thresholding the extrapolated boundary images and varying the threshold parameter. We sampled synthetic billiard sequences using the following parameters.

Table size: Side length randomly sampled from {96,128,160,192,256} pixels.

Ball velocity: Randomly sampled from [{-3,..,3},{-3,..,3}] pixels.

Ball size: Constant, with a radius of 13 pixels.

Initial Position: Uniformly over the table surface.

Real data collection. We captured a novel data-set of real billiard table sequences on a mini-billiard table. Frame rate was set to 120 per second to minimize motion blur. Each sequences consisted of an actor (not visible) striking the ball with a cue stick once. The only motion in the sequences of the dataset are that of the cue stick and the balls. We produce boundary images using the method of [14].

Evaluation on synthetic single ball worlds. We generate a training set using parameters in subsection 4.2. However, to keep our training set as diverse as possible we prefer short sequences. We restrict each sequence to a maximum length of one or two collisions with walls and set a 50% bias of the initial position of the balls being 40 pixels from the walls. We sample 500 such sequences and train the models (from section 3) on these sequences. We then test the models on 30 independent test sequences. We again include the "last input" baseline as a constant extrapolator . We also include a "blind" Convolutional multi-scale Context model (CMSC-BL), which cannot see the table borders. To beat this setup, our models need to learn the physics of ball-wall collisions. We report the results in Table 1.

The Convolutional Multi-Scale Context (CMSC) model performs the best with accurate predictions 20 time-steps into the future – in particular also exceeding the "blind"

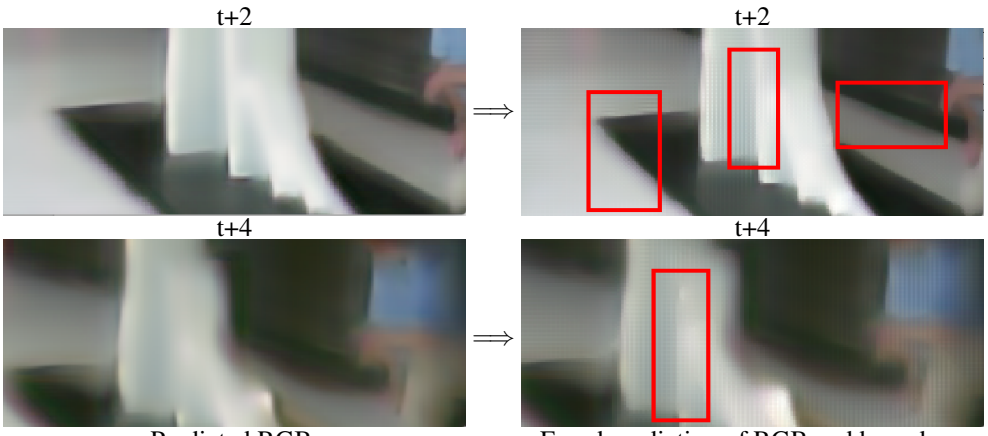

Predicted RGB. Fused prediction of RGB and boundary.

Figure 7: Improving blurry RGB predictions [15] by fusing with sharp boundaries produced by our Convolutional Multi-Scale context model. Red bounding boxes show areas of improvement.

Step	Last Input	C-RNN	CMS	CMSC-BL	CMSC
$t+1$	0.141	0.013	0.282	0.957	**0.987**
$t+5$	0.038	0.006	0.101	0.841	**0.900**
$t+20$	0.002	0.005	0.066	0.347	**0.632**

Table 1: Evaluation of models on single ball billiard table worlds

version (CMSC-BL) that cannot handle bounces. The Convolutional-RNN (C-RNN) architecture (with its bottleneck layer) is unable to learn the physics of the world and produces very blurred output. Without a context, the Convolutional Multi-scale model (CMS) produces inaccurate results at patch borders and thus suffers heavily especially at larger time-steps.

Evaluation on synthetic two and three ball worlds. Worlds with more than one ball also involve ball-ball collisions, which make the physics of such worlds much more complex. To evaluate the models on such worlds we sample 100 and 50 training sequences with two and three balls respectively with a maximum length of 200 frames. We use a curriculum learning approach [1], where we initialize the models with the weights learned on single and two ball worlds respectively. We test the models on 30 independent sequences containing two and three balls respectively. We report the results in Table 2. We also include Convolutional Multi-Scale Context models trained on single ball worlds (CMSC-1B) and two ball worlds (CMSC-2B) in the two and three ball world case respectively. To beat these models learning the physics of ball-ball collisions is necessary. Again, we see accurate extrapolation by the Multi-Scale model even at 20 time-steps in the future. The Convolutional-RNN (C-RNN) architecture performs just as badly as in one ball worlds.

Extrapolation over very long time scales on synthetic data. Although we evaluate only 20 timesteps into the future in Table 1 and Table 2, our models are stable over longer time-horizons. In Figure 8, we show 100 extrapolated frames into the future and visualize them by trails obtained by superposition. We notice a few failure cases where a ball reverse direction mid table and the ball(s) get deformed or disappear altogether.

Evaluation on real billiard sequences. Extrapolation on real billiard table sequences is a challenging test for our models. The table fabric causes rapid deceleration of the ball (compared to the constant velocity in the synthetic sequences). Spin is sometimes inadvertently introduced while striking the ball as well as the segmentation algorithm on the observed frames introduces artifacts. The boundaries are not always consistent across frames of a sequence and they are jagged and change shape. We collect 350 real billiard table sequences, with one ball, as our training set. To deal with deceleration, we experiment with increasing the number of input frames. We train our Convolutional Multi-Scale context model (CMSC-6in) with six input frames. We report the results of evaluation (best F-measure) on 30 independent sequences in Table 3. The low performance of our method compared to the last input is caused by the instability and jaggedness of the static boundaries table and hands. Our method is able to propagate the motion of the ball. We show qualtitive results in Figure 9. The model of [15] is not able to propagate the motion of the ball. We also demonstrate cases as trails, where our model produces stable extrapolations 20 and 50 time-steps into the future.

5. Conclusion

We demonstrate boundary extrapolation that yields accurate and sharp results. Our proposed Convolutional Multi-Scale Context architecture fulfils three key properties for long-term boundary prediction: i) A wide receptive field

	Evaluation on two ball worlds				Evaluation on three ball worlds			
Step	**Last Input**	**C-RNN**	**CMSC-1B**	**CMSC**	**Last Input**	**C-RNN**	**CMSC-2B**	**CMSC**
$t+1$	0.246	0.013	0.966	**0.969**	0.246	0.023	0.967	**0.968**
$t+5$	0.114	0.008	0.848	**0.896**	0.118	0.012	0.890	**0.892**
$t+20$	0.101	0.007	0.612	**0.681**	0.090	0.011	0.664	**0.700**

Table 2: Evaluation of models on complex billiard table worlds

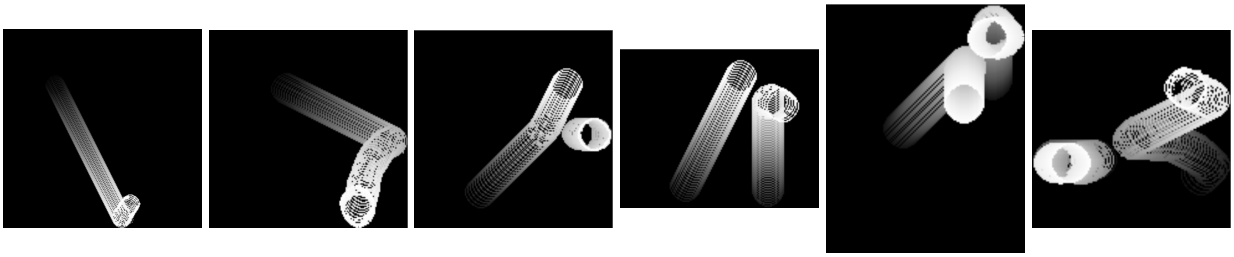

Figure 8: Trails produced by super-imposing extrapolated boundaries.

Groundtruth at t + 5 | RGB prediction [15] | Our Extrapolation.

Trail up to t + 20 | Trail up to t + 50 | Trail up to t + 50

Figure 9: Top row compares RGB prediction of [15] to our extrapolated boundaries. Bottom row shows prediction trails produced by super-imposing extrapolated boundaries.

Step	**Last Input**	**CMSC-6in**
$t+1$	0.253	0.180
$t+5$	0.247	0.174
$t+20$	0.242	0.167

Table 3: Evaluation of models on real ball billiard sequences

which allow the model to learn complex spatio-temporal dependencies. ii) Accurate prediction at each time-step with a fully convolutional setup without any bottleneck layers. iii) The context which allows for the sharing of information thus leading to global consistency. We present long-term prediction results on synthetic scenes involving billiard tables showing that this model can be used to accurately predict the state of such worlds far into the future. This also shows that the model developed an intuitive notion of physics from raw visual input. Moreover, extrapolated boundaries could be used for sharper RGB prediction and could lend itself to formulating expectations over future frames in advanced video segmentation methods. Finally, we have shown first prediction results on real billiard video, that still present many open challenges.

References

[1] Y. Bengio, J. Louradour, R. Collobert, and J. Weston. Curriculum learning. In *ICML*, 2009. 7

[2] J. Chang, D. Wei, and J. Fisher. A video representation using temporal superpixels. In *CVPR*, 2013. 2

[3] E. L. Denton, S. Chintala, A. Szlam, and R. Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In *NIPS*, 2015. 3

[4] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 249–256. ACM, 2002. 5

[5] K. Fragkiadaki, P. Agrawal, S. Levine, and J. Malik. Learning visual predictive models of physics for playing billiards. *arXiv:1511.07404*, 2015. 1, 2

[6] F. Galasso, M. Keuper, T. Brox, and B. Schiele. Spectral graph reduction for efficient image and streaming video segmentation. In *CVPR*, 2014. 2

[7] F. Galasso, N. Nagaraja, T. Cardenas, T. Brox, and B. Schiele. A unified video segmentation benchmark: Annotation, metrics and analysis. In *CVPR*, 2013. 2, 4

[8] V. Jain, J. F. Murray, F. Roth, S. Turaga, V. Zhigulin, K. L. Briggman, M. N. Helmstaedter, W. Denk, and H. S. Seung. Supervised learning of image restoration with convolutional networks. In *CVPR*, 2007. 2

[9] K. Jarrett, K. Kavukcuoglu, M. Ranzato, and Y. LeCun. What is the best multi-stage architecture for object recognition? In *CVPR*, 2009. 3

[10] A. Khoreva, R. Benenson, F. Galasso, M. Hein, and B. Schiele. Improved image boundaries for better video segmentation. In *ECCV Workshop*, 2016. 4

[11] E. P. Krotkov. *Active computer vision by cooperative focus and stereo*. Springer Science & Business Media, 2012. 5

[12] A. Lerer, S. Gross, and R. Fergus. Learning physical intuition of block towers by example. In *ICML*, 2016. 1, 2

[13] W. Li, A. Leonardis, and M. Fritz. Visual stability prediction and its application to manipulation. *arXiv:1609.04861*, 2016. 1, 2

[14] K. Maninis, J. Pont-Tuset, P. Arbeláez, and L. V. Gool. Convolutional oriented boundaries. In *ECCV*, 2016. 6

[15] M. Mathieu, C. Couprie, and Y. LeCun. Deep multi-scale video prediction beyond mean square error. *arXiv:1511.05440*, 2015. 1, 2, 3, 4, 5, 7, 8, 10, 11

[16] V. Michalski, R. Memisevic, and K. Konda. Modeling deep temporal dependencies with recurrent grammar cells. In *NIPS*, 2014. 2

[17] R. Mottaghi, H. Bagherinezhad, M. Rastegari, and A. Farhadi. Newtonian image understanding: Unfolding the dynamics of objects in static images. In *CVPR*, 2016. 2

[18] P. Ochs, J. Malik, and T. Brox. Segmentation of moving objects by long term video analysis. *TPAMI*, 2014. 2

[19] M. Ranzato, A. Szlam, J. Bruna, M. Mathieu, R. Collobert, and S. Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv:1412.6604*, 2014. 2, 3

[20] J. Revaud, P. Weinzaepfel, Z. Harchaoui, and C. Schmid. Epicflow: Edge-preserving interpolation of correspondences for optical flow. In *CVPR*, 2015. 5

[21] N. Srivastava, E. Mansimov, and R. Salakhutdinov. Unsupervised learning of video representations using lstms. In *ICML*, 2015. 2, 3

[22] I. Sutskever, G. E. Hinton, and G. W. Taylor. The recurrent temporal restricted boltzmann machine. In *NIPS*, 2009. 2

[23] J. Wu, I. Yildirim, J. J. Lim, B. Freeman, and J. Tenenbaum. Galileo: Perceiving physical object properties by integrating a physics engine with deep learning. In *NIPS*, 2015. 2

6. Appendix

We include additional image and video results (as files) in this section. We include the following video files (click on the video name for link) with supporting descriptions in this document,

- **vsb100.mp4**, (described in subsection 6.1).
- **singleball_syn.mp4**, (described in subsection 6.2).
- **multipleball_syn.mp4**, (described in subsection 6.2).
- **real_billiards.mp4**, (described in subsection 6.2).

6.1. Example extrapolations on VSB100.

We show extrapolations at one, two and four steps (t + 1, t + 2, t + 4) in the future from a fixed time point on the dominoes and airplane sequences of VSB100. We include all methods except the Convolutional Multi-Scale Context, for which example extrapolations have been shown in Figure 6 of the main article. We show the extrapolations in Figure 10.

As expected from the quantitative performance in Figure 5 of the main article, the "Optic flow" baseline does not perform well. This method incorrectly translates the boundaries which lead to many boundaries being missed especially at t + 4. The Convolutional - RNN model produces very blurry extrapolations which leads to disappearance of boundaries at t + 4. The Multi-Scale model of [15] with Adversarial loss, produces jagged boundaries compared to both of our Convolutional Multi-Scale architectures.

We also include as video, extrapolations using our best performing Convolutional Multi-scale Context model at one, two and four steps (t + 1, t + 2, t + 4) in the future from all frames of the dominoes and airplane sequences of VSB100 (vsb100.mp4). In the video, we see accurate extrapolation in many cases up to four frames into the future (due to mostly smooth motion in these sequences). However, we see incorrect extrapolation in case of non-deterministic motion and high acceleration.

6.2. Example extrapolations over very long time-scales on billiard table sequences.

We include as video example extrapolations 100 time-steps into the future for single and two ball synthetic billiard table sequences (singleball_syn.mp4, multiple-ball_syn.mp4). The Convolutional Multi-scale Context model is given the first four input frames of each sequence. The model then extrapolates the next 100 frames of the sequence. In case of single ball billiard worlds, we see accurate extrapolation in most cases to up to about 75 frames. After that we observe one or more of the following: i) Deformation of the ball. ii) Curvature of the ball trajectory. iii) Change in velocity of the ball. In case of multiple balls (in multipleball_syn.mp4), additionally we observe that after ball-ball collisions, the direction of movement of the balls is correct but speed is usually incorrect. We also include example extrapolations 100 time-steps into the future for real billiard table sequences as video (in real_billiards.mp4). We observe that, our Convolutional Multi-scale Context model is able to extrapolate the motion of the ball well initially even in the presence of rapid deceleration. However, our model is not able to deal with collisions (due to few example collisions in the training data). The table and hand boundaries are distorted with time, due to fluctuations in the initial 6 input frames.

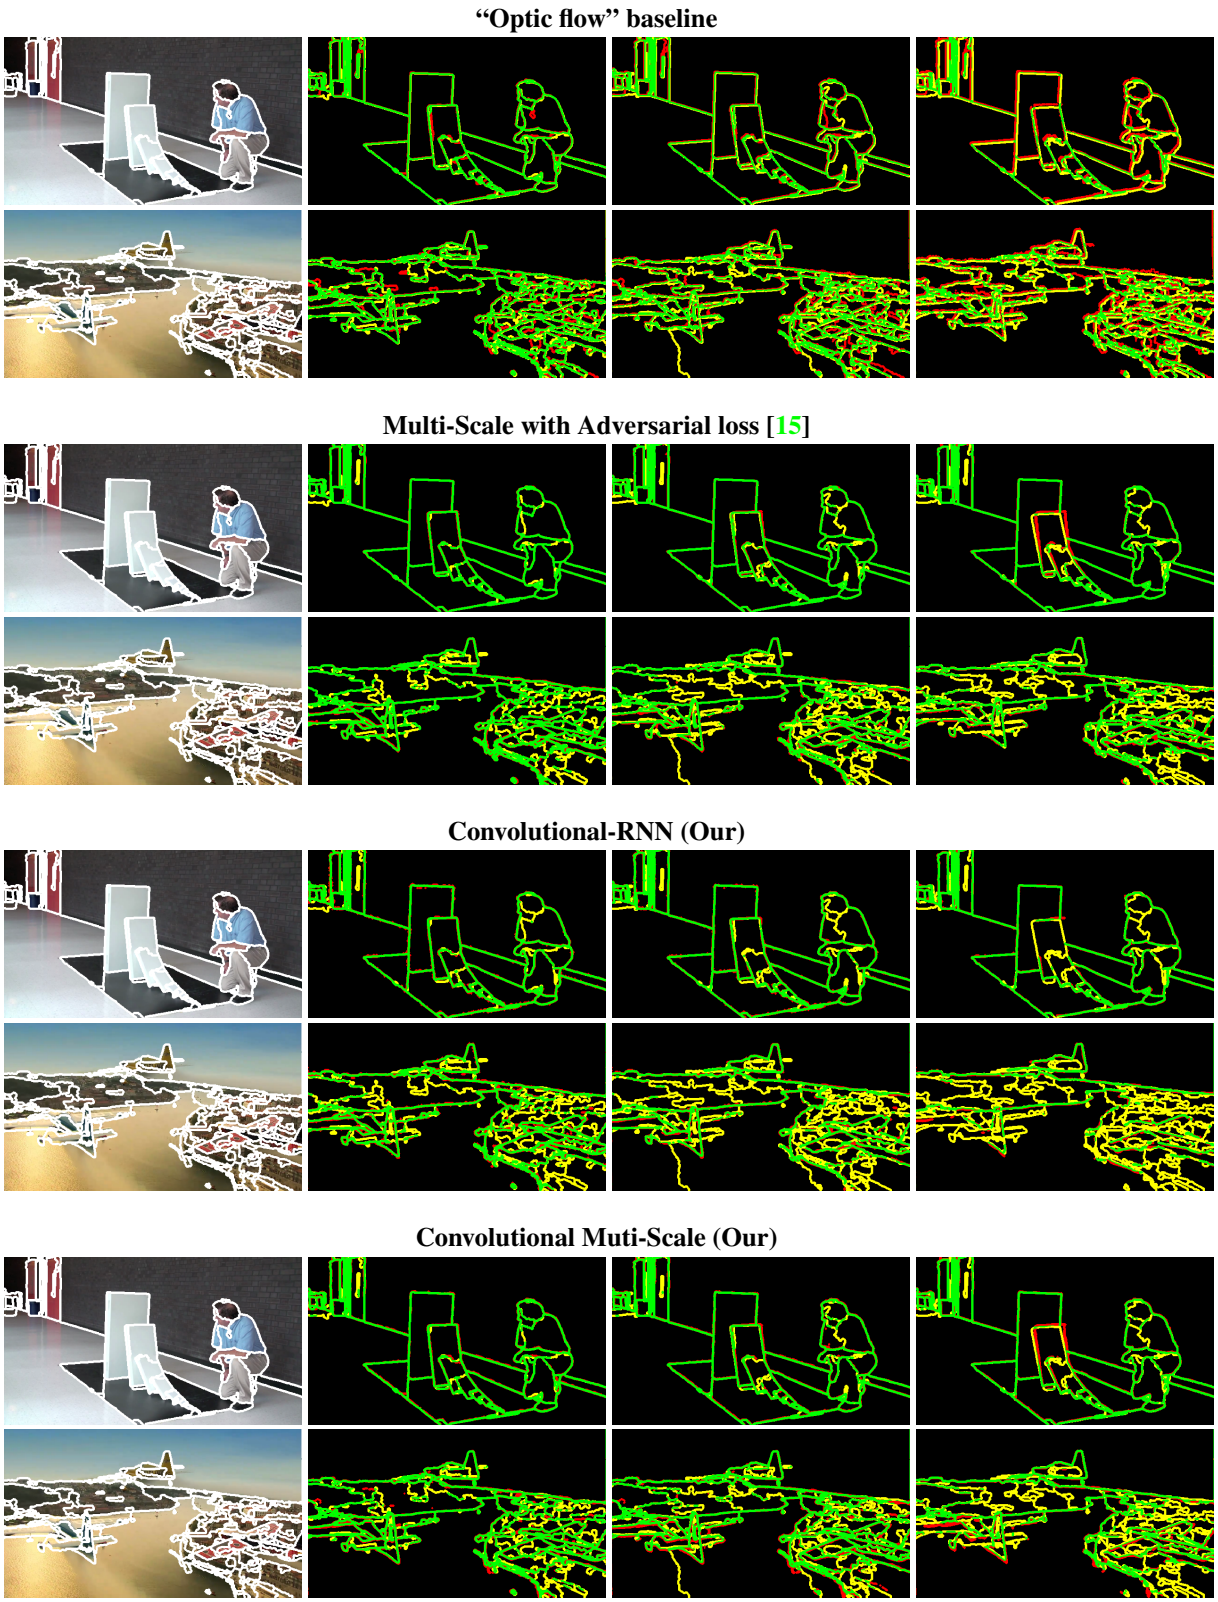

Figure 10: Groups of two rows, Top row in group to bottom in group: Extrapolations on dominoes and airplane sequences from VSB100. Correct boundaries predictions are encoded in green. Missed boundaries are encoded in yellow. Wrong boundaries are encoded in red. Columns left to right: Last frame that the models observe (overlayed with RGB), extrapolation one, two and four time-steps into the future.